

results further demonstrated the equivalent predictive power obtained with the same-subject model, the cross-subject model, and the cross-study model.

[0104] The Clarke EGA is also performed for each of the three studies using the same-subject model predictions (scenario I). The composite result of each analysis is plotted on a separate graph (not shown). Of the 3,600 entries (400 data points \times 9 subjects) for the iSense study, 3,564 points (99.0%) lay in zone A, 35 in zone B, and 1 in zone D. Of the 7,200 entries (400 \times 18) for the Guardian RT study, 7,150 points (99.3%), 32 points, and 18 points lay in zones A, B, and D, respectively. Similarly, of the 2,800 entries of the DexCom study, 2,787 (99.5%), 12, and 1 lay in zones A, B, and D, respectively. These results demonstrated the clinical utility of the predictive models.

[0105] To verify that the employed datasets do not correspond to well-treated diabetic patients with glucose levels mostly within the euglycemic range and that the filtering procedure does not over-smooth the raw data, the number of hypo- and hyperglycemic episodes in the raw, smoothed, and predicted data are calculated. A lower threshold of 3.9 mmol/l (70 mg/dl) and an upper threshold of 10 mmol/l (180 mg/dl) was adopted; and, an inter-episode separation of at least 30 minutes and a minimum of 30 minutes (seven consecutive data points) outside the euglycemic range were required to count the excursion as a hypo- or hyperglycemic episode. FIG. 11 is a table illustrating the cumulative number of hypo- and hyperglycemic episodes and related statistics (averaged over the corresponding subjects) for the raw, smoothed, and predicted data for each of the three studies. The results confirmed that the subjects did exhibit glucose excursions and that the filtering did not significantly smoothed them out. Overall, the models correctly predicted 89 out of 93 hyperglycemic episodes and 20 out of 23 hypoglycemic episodes.

[0106] For instance, for the iSense study, the average minimum glucose levels (in mmol/l) was 3.95, 4.38, and 4.28 for the raw data, smoothed data, and predicted data, respectively. The average maximum glucose levels (in mmol/l) were 15.81, 14.70, and 14.87 for the raw data, smoothed data, and predicted data, respectively. The average mean glucose levels (in mmol/l) were 8.72, 8.72, and 8.69 for the raw data, smoothed data, and predicted data, respectively; and the average standard deviations were 2.61, 2.52, and 2.55 for the raw data, smoothed data, and predicted data, respectively. The total number of hyperglycemic episodes were 25, 24, and 24 for the raw data, smoothed data, and predicted data, respectively; and, the total number of hypoglycemic episodes were 4, 3, and 3 for the raw data, smoothed data, and predicted data, respectively.

[0107] The portability properties demonstrated by the models herein are attributed to two factors: the conserved nature of the frequency content in the glucose signal of diabetic patients and the properties of the modeling approach. The dynamics in the blood glucose time-series signal of diabetic patients can be characterized by four distinct frequency ranges. These different frequency ranges characterize different physiologic mechanisms and are best described by the periodicity of their oscillations. The highest frequency range, with periods between 5 and 15 minutes, is generated by pulsatile secretion of insulin. The second highest, ultradian glucose oscillations, corresponds to periods between 60 and 120 minutes. Exogenous inputs, such as meals and insulin, generate oscillations with periods between 150 and 500 min-

utes; and, finally, circadian oscillations are responsible for the low-frequency range, with periods longer than 700 minutes.

[0108] Analysis of the time-series glucose signals of all subjects in the three studies supports these findings and shows that the frequency content in the signals is conserved across subjects. FIG. 12 is a graph illustrating the power spectrum density profiles for each of the three studies, averaged over the subjects in each study. While the amplitudes of the profiles are different for each of the studies, the periodicity (i.e., the location of the peaks on the x-axis) is conserved across the studies. The conservation of biological rhythms, such as the circadian rhythm, across species, or even kingdoms, is a known phenomenon.

[0109] This similarity in the frequency content of the glucose signals is exploited by the predictive AR models herein. Periodic signals, like glucose concentration, are characterized by three parameters: amplitude, frequency, and phase of the underlying oscillations. However, a property of AR models is their invariance with respect to a signal's amplitude and phase, and sole dependency on its frequency. The sequence of the AR model coefficients captures and represents the frequency content of a time-series signal. Therefore, the development of the predictive AR models from signals with similar frequency content produced similar (or portable) models, regardless that different time-series signals recorded from different subjects had different amplitudes and initial phases. This invariance of the AR model coefficients to the glucose signal's amplitude and phase affords model portability across subjects with type 1 and type 2 diabetes. Type 1 diabetes patients usually have larger glucose-level variations than type 2 patients. However, if these variations contain the same frequency information, the predictive AR models herein are portable across them. Moreover, because of the frequency-dependent nature of the AR model coefficients, information concerning exogenous inputs, such as meals and exercise, is automatically incorporated into the models if this information is present in the training data.

[0110] However, if some of the subjects from the training data are nondiabetic and fasting, the models' portability could be jeopardized because the glucose dynamics are different in this case. This is particularly relevant for the highest-frequency component of the glucose time-series signal, i.e., the shortest periods spanning between 5 and 15 minutes, because while these periods are prominent in nondiabetic, fasting individuals, they are absent in diabetic patients. In diabetic patients, insulin-generating cells responsible for pulsatile secretion of insulin are severely handicapped, essentially eliminating the 5-15 minute periods from the glucose signals. Moreover, the blood-to-interstitial transport acts as a low-pass filter, reducing the high-frequency dynamics in the CGM signals, which are further attenuated by the filtering procedure utilized herein.

[0111] The filtering procedure, used to attenuate any remaining high-frequency component in the signal to yield consistent AR coefficients and robust models, does not significantly impact the ability to capture hypo- and hyperglycemic episodes; and hence, the clinical usefulness of at least one embodiment of the invention. FIG. 11 shows that the predictive models herein correctly predicted 96% of the hyperglycemic episodes and 87% of the hypoglycemic episodes present in the three studies.

[0112] Another contributing property for the predictive AR model portability relates to the limits imposed on the model coefficients by the constrained least squares method. Besides